Uncoupled analysis of stochastic reaction networks in fluctuating environments

Christoph Zechner[1] and Heinz Koeppl[1,2]

[1]Automatic Control Laboratory, ETH Zurich, Switzerland
[2]Technische Universität Darmstadt, Darmstadt, Germany

arXiv:1405.0576v1 [q-bio.QM] 3 May 2014

The dynamics of stochastic reaction networks within cells are inevitably modulated by factors considered extrinsic to the network such as for instance the fluctuations in ribsome copy numbers for a gene regulatory network. While several recent studies demonstrate the importance of accounting for such extrinsic components, the resulting models are typically hard to analyze. In this work we develop a general mathematical framework that allows to uncouple the network from its dynamic environment by incorporating only the environment's effect onto the network into a new model. More technically, we show how such fluctuating extrinsic components (e.g., chemical species) can be marginalized in order to obtain this decoupled model. We derive its corresponding process- and master equations and show how stochastic simulations can be performed. Using several case studies, we demonstrate the significance of the approach. For instance, we exemplarily formulate and solve a marginal master equation describing the protein translation and degradation in a fluctuating environment.

Introduction

Biochemical systems involving low-copy molecules demand for mathematical models that account for the intrinsic stochasticity [1]. In recent years, however, realization has grown that intrinsic noise alone cannot account for the observed substantial phenotypic variability among isogenic cells. That is, fluctuations in the intracellular environment, commonly termed extrinsic noise, represent an additional source of variability [2, 3, 4].

Several recent studies focus on separating intrinsic and extrinsic fluctuations through dual-reporter measurements [5, 6]. Other approaches model extrinsic noise through certain parameters (i.e., the translation rate) of a kinetic model which is calibrated subsequently using flow-cytometry [7, 8] or time-lapse microscopy data [9, 10]. All of those approaches have in common that they consider the biochemical process under study – i.e., the expression of a gene – as a small subpart that is embedded into a larger dynamical system. Accordingly, they rely on augmenting the original kinetic model by certain environmental components which are assumed to be fixed but random [11, 10, 8, 9] or fluctuating over time [5, 12]. In fact, such models agree very well with the variability that is observed experimentally, but on their downside, suffer from the increased dimensionality - somehow defeating the original purpose of tractable dedicated models.

A natural question arising in that context is whether we can find a proper dynamical description of just the system of interest as if it was still embedded into its stochastically modulating environment. In other words, we aim to find a "self-contained" stochastic model that summarizes all system behaviors attainable under all possible realizations of the extrinsic fluctuations. Such models could then be used to perform an uncoupled analysis of a reaction network subject to extrinsic noise. The mathematical correct answer that we provide in this work is the marginalization of the system dynamics with respect to those extrinsic fluctuations. Interestingly it turns out that the resulting model exploits its own stochasticity to emulate the effect of extrinsic noise, leading to a self-exciting process. A simple instance of such self-excitation is the Polya urn scheme[1], which is known to be equivalent to Bernoulli trials marginalized over random (and here correspondingly extrinsic) success rates [13]. Intuitively, due to its self-excitation, the number of draws of the same color for a Polya urn over repeated draws displays a much richer and dispersed dynamics than the number of sucessful draws in a Bernoulli trial with fixed success rate.

For the purpose of inference we recently proposed a first attempt of such marginalization for the special case of fixed but random environmental conditions [9]. In this work we develop a general mathematical framework from which the uncoupled dynamics can be constructed in a principled manner, regardless whether the environment is constant or dynamically changing.

Results

Mathematical modeling We describe the time-evolution of a stochastic reaction network by a continous-time Markov chain (CTMC) X with M chemical species and N reaction channels. The system state at time t is denoted $X(t)$ and we write its random[2] path on time intervals $[0, a]$ as $\mathbf{X}_a$. Furthermore, we assume that X depends on another multivariate Markov process Z through its hazard functions in the form

$$h_i(x, z) = c_i(z)g_i(x), \tag{1}$$

with c_i some positive function and g_i a polynomial determined by the law of mass-action, for instance. For

[1]At each draw from the Polya urn with balls of two colors the drawn ball and a fixed number of new balls of the same color as the draw are placed in the urn.

[2]Throughout we will follow the usual convention to refer to upper-case and lower-case versions of a symbol as a random variable and its realization, respectively.

reactions independent of Z, we thus have $c_i(z) \equiv c_i$. Typically, Z is another jump or diffusion process corresponding to a set of modulating *environmental* species or conditions that are considered extrinsic to the system of interest, whereas the species in X represent the actual system of interest. For example, Z could be the fluctating ribosome copy numbers affecting the kinetics of a gene regulatory network represented by X. Although a more general treatment is possible, we assume a feed-forward structure between Z and X, which means that Z modulates X but not vice-versa. Consequently, the dynamics of the joint system $Y(t) = (Z(t), X(t))$ can be described by a marginal Markov process Z together with a conditional Markov chain $X \mid Z$.

Uncoupled dynamics Mathematical descriptions of the joint system $Y(t)$ are readily obtained using available techniques for modeling Markovian dynamics [5, 14, 12]. For complexity reasons, however, we aim for models that can properly describe *only* the interesting components $X(t)$. In order to see that marginalization over Z yields the desired model, let us first consider two dependent random variables A and B described by a joint probability distribution $p(a, b) = p(a \mid b)p(b)$. If we are interested in analyzing A under all possible values of b, we need to average the probability at $A = a$ over all possible values of b, i.e.,

$$p(a) = \int p(a, b)\mathrm{d}b = \mathbb{E}\left[p(a \mid B)\right].$$

Note that as a consequence of averaging probabilities, any value a possible (i.e. has non-zero measure) under the joint $p(a, b)$ is possible under the marginal $p(a)$, while this does not necessarily apply to $p(a \mid b)$ for any choice of b.

In case of the coupled processes $Z(t)$ and $X(t)$, we analogously *marginalize* the joint Markov chain Y with respect to the environmental process Z. While such a marginalization involves several difficulties, the idea remains the same: we try to construct an uncoupled process X which directly admits the marginal path distribution $p(\mathbf{x}_t) = \mathbb{E}\left[p(\mathbf{x}_t \mid \mathbf{Z}_t)\right]$, bypassing the intractable averaging over all possible extrinsic histories. As a result, we obtain a jump process which - in contrast to the conditional process $X \mid Z$ - no longer depends on the environmental species in Z. We remark that a straightforward marginalization of the joint master equation of Z and X generally leads to intractable propensities [15, 5]. Based on the innovation theorem [16] we demonstrate in section S.1 in the SI Appendix that the hazard functions of the uncoupled process can be generally written as

$$\lambda_i(\mathbf{X}_t) = \mathbb{E}\left[c_i\left(Z(t)\right) \mid \mathbf{X}_t\right] g_i(X(t)), \tag{2}$$

where the expectation is taken with respect to the conditional distribution $\pi\left(z, t \mid \mathbf{x}_t\right)$. The latter describes the conditional probability of the environmental process $Z(t)$ given the entire history of process X until time t[3]. Using the expected value of that distribution, the feed-forward influence of Z on the hazard functions of X can be replaced by a deterministic function of X, which no longer depends on the actual state of Z. Instead, the marginal process X becomes *self-exciting*, meaning that it exerts a feedback on itself. Note that the uncoupled process X is no longer Markovian, since the conditional expectation - and hence the hazard functions - depend on the full process history $\mathbf{X}_t$. A schematic illustration of that uncoupling is given in Fig.1.

[3]More precisely we would need to say that $\mathbf{X}_t$ is a filtration of X.

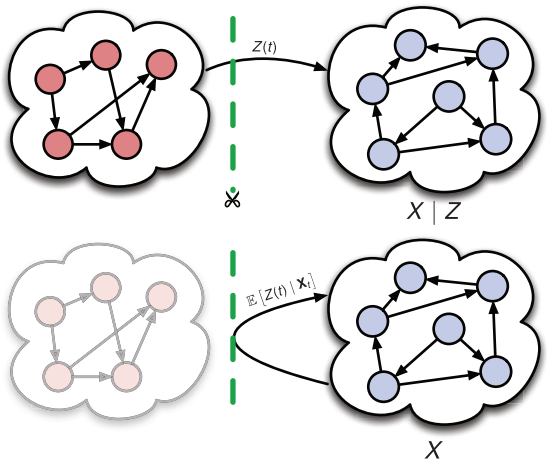

Figure 1: Uncoupled stochastic dynamics. The environmental process Z modulates the dynamics of the process under study X, e.g., through one of its hazard functions. Marginalization with respect to Z yields the uncoupled dynamics of X, whereas the original dependency on the environment Z is replaced by its optimal estimator given the history of X. Consequently, the marginal process X is self-exciting, i.e., it exerts a feedback on itself.

Solving the accompanying filtering problem Although the construction of the uncoupled dynamics is general, any practical implementation thereof will depend on an explicit computation of the conditional expectation in Eq. 2. This expectation estimates the environmental state $Z(t)$ given the full history of the uncoupled process $\mathbf{X}_t$ and therefore, can be understood as the solution to a *stochastic filtering* problem [17]. Filtering techniques deal with the problem of optimally reconstructing a hidden stochastic process at time t from noisy observations of that process up to time t. In the situation considered here, the hidden process corresponds to the environment $Z(t)$, which gets reconstructed from the "observed" history $\mathbf{X}_t$ through the conditional mean in Eq. 2.

We assume that the environment $Z(t)$ admits a probability distribution $p(z, t)$ described by a Kolmogorov-forward equation of the form

$$\frac{\partial}{\partial t}p(z, t) = \mathcal{A}p(z, t), \tag{3}$$

where $\mathcal{A}$ represents the temporal change of $p(z, t)$, i.e., is the infinitesimal generator of Z. For instance, if Z is a diffusion process, $\mathcal{A}$ corresponds to the Fokker-Planck operator, while in case of a CTMC, $\mathcal{A}$ is given by the difference operator of the chemical master equation (CME). In terms of filtering, Eq. 3 corresponds to the process model of Z. Furthermore, we know that at a given time t, the solution of X can be written as a sum of independent but time-transformed Poisson processes [18], each

of them corresponding to a particular reaction channel. Consequently, the observation model is given by a set of Poisson counting observations with the hazard functions given in Eq.1. This is closely related to Markov-modulated Poisson processes [19] and their corresponding optimal filtering [20].

While a more general treatment is provided in the SI Appendix, we assume in the following that a one-dimensional process Z is modulating X through its k-th reaction of order zero. We further restrict ourselves to the case where c_k is a linear function of z, i.e., $c_k(z) = c_k z$. Under those assumptions, it can be shown that the conditional process $Z(t) \mid \mathbf{X}_t$ follows a filtering distribution $\pi(z, t \mid \mathbf{x}_t) \equiv \xi(t)\tilde{\pi}(z, t)$ with

$$\mathrm{d}\tilde{\pi}(z,t) = \left[\mathcal{A}\tilde{\pi}(z,t) - c_k z\tilde{\pi}(z,t)\right]\mathrm{d}t + \left[z-1\right]\tilde{\pi}(z,t)\mathrm{d}R_k(t), \tag{4}$$

with $\xi(t)$ a time-dependent normalizing factor independent of z and $R_k(t)$ the number of reactions of type k up to time t in $\mathbf{X}_t$. Thus, Eq.4 describes a scaled version of the normalized filtering distribution. The latter shows an implicit dependency on its own mean (see Methods and section S.2 in the SI Appendix) and is therefore complicated to handle numerically. In contrast, once we have numerically solved for $\tilde{\pi}$, it can be easily rescaled such that it integrates (or sums up) to one for all t. Note that Eq. 4 is a stochastic partial differential equation (SPDE) in case Z describes a diffusion process or a stochastic difference-differential equation (SDDE) if Z is a CTMC. In the latter case, the solution of Eq.4 can be compactly written as

$$\tilde{\Pi}(t) = e^{(Q-c_k\Lambda)t}\Lambda^{R_k(t)}\Pi_0, \tag{5}$$

with $\tilde{\Pi}(t) = (\tilde{\pi}(0,t), \ldots, \tilde{\pi}(L-1,t))^T$, L the number of reachable states of Z, $Q \in \mathbb{R}^{L\times L}$ the generator matrix of Z, $\Lambda = \mathrm{diag}(0, \ldots, L-1)$ and $\Pi_0 \in \mathbb{R}^L$ the initial distribution over Z.

In order to evaluate Eq. 2, we only require the mean (i.e., the first moment) of the filtering distribution, i.e., $M_1(t) = \mathbb{E}\left[Z(t) \mid \mathbf{X}_t\right]$. In general, however, the mean also depends on the second-order moment, which in turn depends on the third-order moment and so forth. We show in the Methods section that the (non-central) filtering moment dynamics up to order i can be generally written as

$$\begin{aligned}
\mathrm{d}M_1(t) &= \left[\mathcal{D}_1(t) - c_k(M_2(t) - M_1(t)M_1(t))\right]\mathrm{d}t \\
&\quad + \frac{M_2(t) - M_1(t)M_1(t)}{M_1(t)}\mathrm{d}R_k(t) \\
&\vdots \\
\mathrm{d}M_i(t) &= \left[\mathcal{D}_i(t) - c_k(M_{i+1}(t) - M_1(t)M_i(t))\right]\mathrm{d}t \\
&\quad + \frac{M_{i+1}(t) - M_1(t)M_i(t)}{M_1(t)}\mathrm{d}R_k(t),
\end{aligned} \tag{6}$$

where $\mathcal{D}_j(t)$ refers to the prior dynamics of the j-th moment. Although Eq.6 is generally infinite-dimensional, there are several relevant scenarios, for which the moment dynamics are *closed*, i.e., only depend on higher-order moments up to a certain order. This is for instance the case, if $Z(t)$ is a Cox-Ingersoll-Ross process or any finite state Markov chain. On the other hand, if the moment dynamics are infinite-dimensional, suitable assumptions on the filtering distribution π can be imposed to yield a closed moment-dynamics (see S.3 in the SI Appendix). An important closure is found by analyzing Eq. 5: especially for large c_k, the conditional distribution of Z is predominantly driven by the term $e^{-c_k\Lambda t}\Lambda^{R_k(t)}$, suggesting that it can be well approximated by a Gamma-distribution. We note that the Gamma-distribution is fully characterized by two parameters – or equivalently – its first two moments $M_1(t)$ and $M_2(t)$. As a consequence, we may express the third order moment as a function of the first two moments, i.e., $M_3(t) = -M_1(t)M_2(t) + 2M_2^2(t)/M_1(t)$, such that the second conditional moment closes as

$$\begin{aligned}
\mathrm{d}M_2(t) = &\left[\mathcal{D}_2(t) - 2c_k\frac{M_2(t)}{M_1(t)}\left(M_2(t) - M_1^2(t)\right)\right]\mathrm{d}t \\
&+ 2\left[\frac{M_2^2(t)}{M_1^2(t)} - M_2(t)\right]\mathrm{d}R_k(t).
\end{aligned} \tag{7}$$

Further discussion on this closure is provided in section S.3 in the SI Appendix.

Stochastic simulation Although the uncoupled dynamics of X are non-Markovian, the Markov property can be enforced by virtually extending the state space by the filtering mean of Eq. 6, which summarize the history of X. As a result, one can simulate sample paths of the uncoupled process using standard methods that can account for the explicit time-dependency of the hazard functions [21]. In general, such algorithms rely on the generation of random waiting times for each of the reaction channels. All reactions that are independent of $Z(t)$ will retain their exponentially distributed waiting times. In contrast, the time τ_k that passes until a reaction of type k happens is distributed according to

$$P_k(\tau_k < s \mid \mathbf{X}_t) = 1 - e^{-c_k\int_0^s M_1(t+T)\mathrm{d}T}. \tag{8}$$

We note that as long as no reaction of type k happens, $\mathrm{d}R_k(t)$ is zero and hence, $M_1(t)$ is found by solving a set ordinary differential equations (ODEs). Since that solution is not generally known in closed form, we cannot directly sample from Eq. 8. However, several efficient solutions to that problem have been developed in the context of inhomogeneous Poisson processes, e.g., such as the method of *thinning* [22] (see Methods). Once a reaction has fired, the filtering moments need to be updated by the terms multiplying the firing process $\mathrm{d}R_k(t)$ in Eq. 6 (i.e., they exhibit a discontinuity).

Evidently, simulation from Eq. 8 comes at higher cost than simulating from an exponential distribution (e.g., such as performed in standard SSA algorithms), since in general, it relies on the numerical integration of an ODE. However, reactions associated with the environmental part no longer need to be simulated, which yields a significant reduction in computational effort as soon as the environmental network is large and expensive to simulate due to high propensity reactions, for instance.

Fluctuations on different timescales The impact of environmental fluctuations on a dynamical system of interest is as diverse as the timescale on which they operate. For instance, extrinsic noise in the context of gene expression might be slowly varying (e.g., correlates well with the

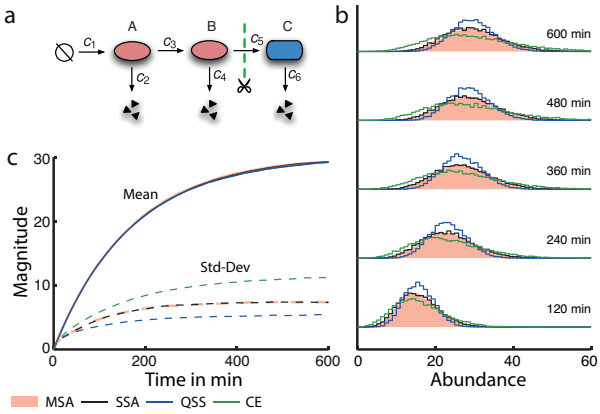

Figure 2: Marginal simulation algorithm. (a) Simple three-stage model. Species A, B and C are modeled as coupled linear birth-death processes, where the coupling is realized by linearly modulating the birth rates of B and C (rate constants $c_1 = 0.003, c_2 = 0.001, c_3 = 0.05, c_4 = 5e-4, c_5 = 1e-5, c_6 = 1e-4$). The uncoupled marginal dynamics of C are obtained by integrating over fluctuations of species A and B (7000 sample paths were used). (b, c) Evaluation of the marginal simulation algorithm. Simulations based on the QSS-approximation neglect a significant portion of variability as opposed to assuming a constant environment (CE) in which case the variability is overestimated. In contrast, the uncoupled dynamics correctly predict the fluctuations on the protein level, while yielding a reduction in computational effort when compared to standard SSA (20min simulation time instead of 46min); correspondingly higher speedup can be achieved for a larger time-scale separation of processes (A,B) versus C.

cell-cycle [23, 24]), while fluctuations in transcription factor abundance might be significantly faster than the expression kinetics downstream. From a technical point of view, timescales range from constant environmental conditions that are random but fixed [25] to regimes where the fluctuations are very fast, such that quasi-steady-state (QSS) assumptions become applicable [15]. A QSS-based approach for simulating a system X in the presence of extrinsic noise Z corresponds to simulating the conditional CTMC $X \mid Z$, where Z is replaced by the mean of Z. The simulation of the joint system (X, Z) become prohibitive if extrinsic fluctuations are fast, while with Eq. 6 the complexity of the marginal process simulation is invariant with respect to the time-scale of the environment. Alternatively, one may try to replace a fluctuating environment Z through a random but fixed enviroment of same variance but this leads to an overestimation of the process variance in X [5], as discussed in a later section. To investigate the two above simplifying assumptions and compare them to the exact solution obtained via SSA and via the marignal process, we performed a simulation study on a linear three-stage birth-death model given in Fig.2a, where only species C is considered of interest in this case. Accordingly, the uncoupled dynamics of C are obtained by integrating the dynamics over the A and B. The results are shown in Fig.2b and Fig.2c.

Propagation of environmental fluctuations and the effective noise Several recent studies [26, 5, 6, 4] are centered around the separation of different noise contributions in biochemical networks. Typically, the law of total variance is employed to decompose the fluctuations of $X(t)$ into parts that are intrinsic to X and parts that come from Z (i.e., are extrinsic to X). Here we found that performing such an analysis on Z instead of X – in conjunction with our decoupling approach – provides a novel way to study how stochasticity is propagated through biochemical networks. Using the law of total variance, we can decompose the total (or unconditional) variance of $Z(t)$ as

$$\mathrm{Var}\left[Z(t)\right] = \mathbb{E}\left[\mathrm{Var}\left[Z(t) \mid \mathbf{X}_t\right]\right] + \mathrm{Var}\left[\mathbb{E}\left[Z(t) \mid \mathbf{X}_t\right]\right]. \quad (9)$$

The two terms on the r.h.s. can be interpreted as follows. Assume we can observe Z only through X. Since X is intrinsically stochastic, a part of the variability of Z is not carried over to X. In Eq. 9, this part (i.e., the *suppressed noise*) corresponds to the first term on the r.h.s. since it quantifies the uncertainty about $Z(t)$ that remains after observing $\mathbf{X}_t$. The second term determines how accurate Z can be reconstructed from trajectories of $\mathbf{X}$. Alternatively, it can be understood as the amount of noise in Z that effectively impacts X (i.e., the *effective noise*). For instance, the environmental process could be characterized by a large variance, but still have only marginal impact on $X(t)$ – depending on the timescale of Z and X.

In order to quantify those terms, we note that the conditional variance within in the first term coincides with the second-order central moment of the filtering distribution from Eq.4. This further implies that it can be computed "on-the-fly" when simulating $X(t)$ using the marginal simulation algorithm which allows an efficient estimation of its expectation. However, in some biologically relevant cases, the effective noise can be determined even analytically, which we demonstrate in the following.

We derive in section S.4 in the SI Appendix that the expected central moments are generally given by

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}\left[M_1(t)\right] &= \mathbb{E}\left[\mathcal{D}_1(t)\right] \\ \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}\left[S_2(t)\right] &= \mathbb{E}\left[\tilde{\mathcal{D}}_2(t)\right] - c_k\mathbb{E}\left[\frac{S_2^2(t)}{M_1(t)}\right]. \end{aligned} \quad (10)$$

The mean in Eq.10 is just the unconditional mean of $Z(t)$, while the derivative of the expected variance shows an additional negative term, causing it to be smaller than the unconditional variance. Let us for instance consider the case where $Z(t)$ follows a Cox-Ingersoll-Ross (CIR) process governed by the SDE

$$\mathrm{d}Z(t) = \theta(\mu - Z(t))\mathrm{d}t + \sigma_Z\sqrt{Z(t)}\mathrm{d}W(t), \quad (11)$$

with θ, μ and σ_Z as real process parameters and $W(t)$ as a standard Wiener process. Note that in this case, Eq. 10 reduces to an autonomous ODE, which for large t yields the relative effective noise $\eta = \mathrm{Var}\left[\mathbb{E}\left[Z(t) \mid \mathbf{X}_t\right]\right] / \mathrm{Var}\left[Z(t)\right]$ at stationarity, i.e.,

$$\eta = 1 + 2\frac{v^2}{c_k}\left(1 - \sqrt{\frac{c_k}{v^2} + 1}\right), \quad (12)$$

where $v = \theta/\sigma_Z$ can be considered a normalized timescale of $Z(t)$ (see section S.4 in the SI Appendix). The computation of the effective noise and its dependency on the environmental timescale is illustrated in Fig. 3.

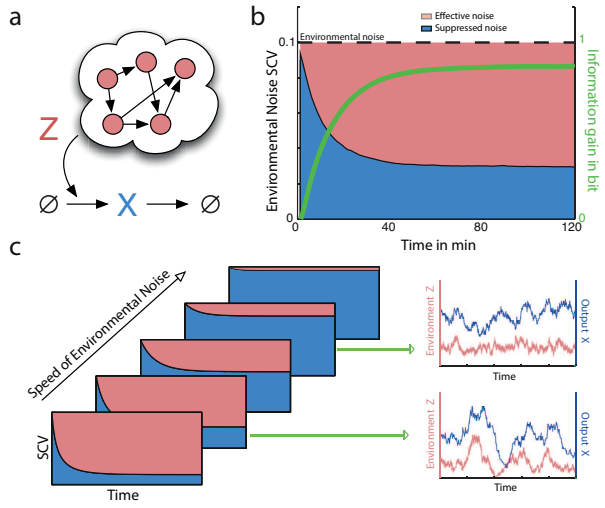

Figure 3: Propagation and suppression of environmental fluctuations. (a) Linear birth-death process in a fluctuation environment. The birth-rate is assumed to be linearly modulated by an environmental stochastic process Z. (b) Calculation of suppressed and effective noise. Individual components were computed analytically by solving an ordinary differential equation (see main text). For orientation, we also show the information gain between $Z(t)$ and $Z(t) \mid \mathbf{X}_t$, computed using the marginal simulation algorithm (green); it can be understood as the gain in information about Z through observing X and it exhibits a monotone relationship with the effective noise. (c) Relation between the effective noise and the speed of the environmental fluctuations. Noise contributions were computed by numerically solving the ODE from Eq.10 for different values of θ (i.e., timescales).

The slow noise approximation (SNA) The effective noise can be understood as a measure of how strong Z impacts X. Only in the special case of a very slow or constant environment, i.e., $\mathcal{D}_i \approx 0$, we see from Eq.10 that for large t, $\mathrm{Var}\left[\mathbb{E}\left[Z(t) \mid \mathbf{X}_t\right]\right] \to \mathrm{Var}\left[Z(t)\right]$, i.e., all variability in Z is transferred to X. Hence, a more noisy but fluctuating environment may induce a similar (or even the same) effective noise in X than a random but fixed environment of the same variance. Consequently, when looking at only snapshot data for X one can generally not infer whether the environment is constant or fluctuating. On the other hand, this implies that we may well approximate the impact of a complicated and dynamically changing environment by a simple random variable of appropriate variance. More specifically, we demand for an equivalent constant environment $\bar{Z}$ such that $\mathrm{Var}\left[\bar{Z}\right] \equiv \mathrm{Var}\left[\mathbb{E}\left[Z(t) \mid \mathbf{X}_t\right]\right]$, where $\mathrm{Var}\left[\mathbb{E}\left[Z(t) \mid \mathbf{X}_t\right]\right] = \sigma^2$ is the effective noise of the original, fluctuating environment Z at stationarity. Let us again consider the birth-death process of Fig. 3a and set the birth rate to one such that any scaling is subsumed in the environmental process Z. With $X_0 = 0$, the abundance of the birth death process at any time is given by $X(t) = R_b(t) - R_d(t)$ with $R_b(t)$ and $R_d(t)$ as counting processes for the birth and death reaction, respectively. We show in section S.5.1 in the SI Appendix that the marginal birth hazard is approximately given by

$$\lambda_b(\mathbf{X}_t) = \lambda_b(R_b(t), t) \approx \frac{\mu^2 + \sigma^2 R_b(t)}{\mu + \sigma^2 t}, \tag{13}$$

with $\mu = \mathbb{E}\left[Z(t)\right]$ the unconditional mean and σ^2 the effective noise of Z, whereas the expression becomes exact for constant and infinitely fast environments. Note that the marginal hazard does not depend on the full history, but only the number of birth-reactions $R_b(t)$ up to time t[4]. In relation to QSS, which assumes that no fluctuations of Z are propagated to X, the found equivalent constant environment with the proper effective noise provides a better approximation for a decoupled simulation of environment and process of interest than QSS.

Using the effective noise, we now aim to find a master equation, which describes the time-evolution of the marginal probability distribution $P(x,t)$. Since λ_b depends on $R_b(t)$ rather than $X(t)$, it appears natural to formulate the master equation in $R_b(t)$ and $R_d(t)$ as well. We remark that since the uncoupled dynamics are non-Markovian, they do not satisfy a conventional master equation. Instead, such processes are described by *generalized master equations* (GME) that can account for memory effects in the dynamics (see S.5 in the SI Appendix for a general derivation and discussion). For the example considered here, one can show that the probability distribution $P(r_b, r_d, t)$ satisfies a GME of the form

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}P(r_b,r_d,t) &= \frac{\mu^2+\sigma^2(r_b-1)}{\mu+\sigma^2 t}P(r_b-1,r_d,t)\\ &+ c_d\left[r_b - r_d + 1\right]P(r_b, r_d-1, t)\\ &- \left(\frac{\mu^2+\sigma^2 r_b}{\mu+\sigma^2 t} + c_d\left[r_b - r_d\right]\right)P(r_b,r_d,t),\end{aligned} \tag{14}$$

that can be solved analytically using generating functions (see S.5.1 in the SI Appendix). From $P(r_b, r_d, t)$ we compute the distribution of X as

$$\begin{aligned}P(x,t) &= \sum_{r_b=x}^{\infty} P(r_b, r_b - x, t)\\ &= \mathcal{NB}\left(x; \frac{\mu^2}{\sigma^2}, \frac{c_d \mu e^{c_d t}}{c_d \mu e^{c_d t} + \left(e^{c_d t} - 1\right)\sigma^2}\right),\end{aligned} \tag{15}$$

i.e., a negative binomial distribution. Eq. 15 provides a surprisingly simple approximate solution for the transient probability distribution of birth death processes in a fluctuating environment. In order to check its validity, we compared the analytical approximate distributions to the ones obtained through SSA for a gene expression model, where the environmental fluctuations are assumed to be due to the mRNA dynamics (see Fig. 4). More specifically, we computed the Kolmogorov distance between the resulting protein distributions as a function the environmental timescale. Apart from the exact correspondence for the limiting time-scales, Fig. 4 indicates that the SNA provides a good approximation regardless of the environmental timescale.

Discussion

There is increasing evidence that models of biochemical networks need to account for both intrinsic and extrinsic noise caused by variations in the intracellular environment. In recent studies, this is done by extending a model's state space by certain environmental species,

[4] That is, $R_b(t)$ is a sufficient statistic for evaluating the conditional expectation $\mathbb{E}\left[Z(t) \mid \mathbf{X}_t\right]$.

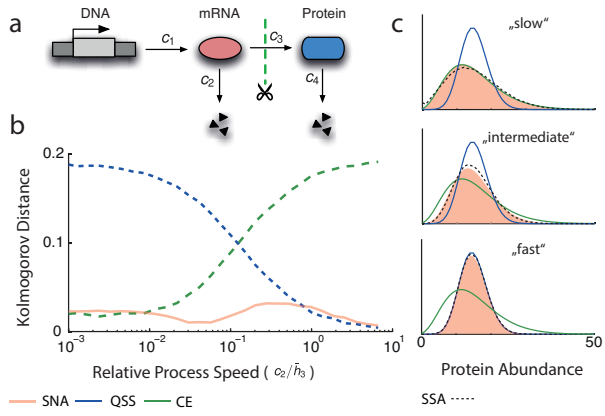

Figure 4: Analytical protein distributions through the slow noise approximation. (a) Two-stage gene expression model. Transcription and translation are modeled through mass-action kinetics with reaction rate constants c_1-c_4. Fluctuations on the mRNA are considered environmental and hence, integrated out in order to obtain a one-dimensional stochastic process describing only the protein. (b) Accuracy of the slow noise approximation. The SNA was compared to the QSS- and CE-approximations by means of the Kolmogorov distance between the respective approximate and exact distribution (SSA) as a function of the relative speed of the mRNA fluctuations, i.e., $c_2/\bar{h}_3$ and $\bar{h}_3 = c_3c_1/c_2$. QSS- and CE approximations break down for slow or fast environmental fluctuations respectively, whereas the SNA yields accurate distributions regardless of the mRNA's timescale. (c) Exemplary distributions obtained through the different approaches in three different regimes (slow, intermediate, fast).

whose dynamics are described along with the actual system of interest. In particular, the resulting system dynamics are described and studied *conditional* on a particular history of the environment and thus, do not provide a coherent description of a dynamical system subject to extrinsic noise. In this work, we derived and analyzed a novel process framework, which is able to describe just the system of interest as if it was still embedded into its environment. In that sense, it permits a mathematically exact way to analyze small parts of networks in an *uncoupled* fashion.

Several recent studies rely on the extreme assumptions that the environmental fluctuations are either infinitely fast or slow. While both strategies may in fact lead to strongly simplified and tractable models, they are characterized by significant approximation errors when considering intermediate environmental timescales (see e.g., Fig. 4b). The approach proposed here allows to uncouple a reaction network from its surrounding environment regardless of the latter's timescale. In that sense, the approach is fully general although practical implementations may rely on efficient but approximate solutions of the discussed filtering problem.

In the context of Monte Carlo simulation the decoupled process can yield a significant reduction in computational effort when compared to standard SSA – especially if the environmental network is costly to simulate. This highlights the role of the provided framework as a general tool to split stochastic biochemical networks into individual parts that are easier to simulate. We believe that it will aid in turning stochastic modeling and simulation techniques more large-scale and more faithful to *in vivo* conditions, where significant environmental fluctuations are present. Moreover, the framework can be used in the model-based design of novel circuit motifs in synthetic biology and is related to the notion of retroactivity [27].

We further demonstrate that the uncoupled dynamics provide a novel analytical tool to study how environmental stochasticity is propagated along coupled reaction networks. For instance, we have shown that the total environmental noise splits up into two terms: one corresponding to the noise that is suppressed and a second term that quantifies the effective noise that is *sensed* by the target network.

In [28] the authors derive a lower bound on a network's ability to suppress fluctuations and show its immediate relation to the uncertainty at which those fluctuations can be estimated – similar to what we defined as effective noise. The methods proposed here allow to not only bound, but fully determine both the suppressed and effective environmental noise. Our results further indicate that two environments with very different timescales may impact a network in a similar way. For instance a fixed but random environment may yield the same effective noise as a fluctuating environment with larger variance. Along those lines, we derived a simple but widely applicable approximation of the transient probability distribution for birth death processes subject to environmental noise. It is based on the idea to approximate a fluctuating environmental process by a simple random variable that impacts the birth death process in an equivalent way. In order to solve for the transient probability distribution we derived a novel generalized master equation for this non-Markovian process.

Methods

Normalized filtering distribution. The unnormalized filtering distribution from Eq.4 does not sum up or integrate to one and therefore, cannot be used to derive statistics such as moments and so forth. We show in section S.2 in the SI Appendix that the normalized filtering distribution is given by

$$\begin{aligned} \mathrm{d}\pi(z,t) = &\Big(\mathcal{A}\pi(z,t) - c_k\left[z - M_1(t)\right]\pi(z,t)\Big)\mathrm{d}t \\ &+ \left[\frac{z - M_1(t)}{M_1(t)}\right]\pi(z,t)\mathrm{d}R_k(t), \end{aligned} \tag{16}$$

which – due to a dependency on the mean $M_1(t)$ – is difficult to integrate numerically. On the other hand, moment dynamics are straight-forward to derive such as described in the following section.

Conditional moment dynamics. The i-th order noncentral moment is computed by multiplying both sides of Eq.16 with z^i and summing (or integrating) over all $z \in \mathcal{Z}$, i.e.,

$$\begin{aligned} \sum_{z\in\mathcal{Z}} z^i \mathrm{d}\pi(z,t) &= \sum_{z\in\mathcal{Z}} z^i \Big(\mathcal{A}\pi(z,t) - c_k\left[z - M_1(t)\right]\pi(z,t)\Big)\mathrm{d}t \\ &\quad + \sum_{z\in\mathcal{Z}} z^i \left[\frac{z - M_1(t)}{M_1(t)}\right]\pi(z,t)\mathrm{d}R_k(t) \\ &= \left[\mathcal{D}_i(t) - c_k(M_{i+1}(t) - M_1(t)M_i(t))\right]\mathrm{d}t \\ &\quad + \frac{M_{i+1}(t) - M_1(t)M_i(t)}{M_1(t)}\mathrm{d}R_k(t), \end{aligned} \tag{17}$$

with $\mathcal{D}_i(t)=\sum_{z\in\mathcal{Z}} z^i \mathcal{A}\pi(z,t)$. The computation of moments in case of multivariate environments is obtained analogously.

Marginal simulation algorithm. As indicated in the main text, the uncoupled dynamics can be simulated using any stochastic simulation algorithm that can deal with time-varying hazard functions. Although more efficient variants might be possible, we make use of a first-reaction method [21] that we combine with a thinning algorithm [22]. The first reaction method is based on drawing waiting times $\tau_1,...,\tau_N$ for each of the N reaction channels and then picking the reaction corresponding to the minimum of those waiting times. First we remark that only the reactions involving $Z(t)$ will be affected by the decoupling scheme and hence, all other reactions will retain their exponential waiting time distributions. In order to simulate from the non-exponential waiting time distribution from Eq.8, we apply a thinning algorithm given by the following steps:

1. Set $\tau=t$.
2. Simulate $\hat{\tau}\sim\text{Exp}(\hat{\lambda})$.
3. Set $\tau=\tau+\hat{\tau}$.
4. Draw $u\sim\mathcal{U}(0,1)$. If $u\leq\frac{M_1(\tau)g_k(\mathbf{x}(\tau))}{\hat{\lambda}}$ return τ. Else, go back to step 2.

Note that the tuning parameter $\hat{\lambda}$ has to be chosen such that $\hat{\lambda}\geq M_1(\tau)g_k(\mathbf{x}(\tau))$ for all $\tau\in[0,T]$, where T is the simulation interval.

Acknowledgements. C.Z. and H.K. acknowledge the support from the Swiss National Science Foundation, grant no. PP00P2 128503 and SystemsX.ch, respectively.

References

[1] McAdams HH, Arkin A (1997) Stochastic mechanisms in gene expression. *Proc Natl Acad Sci USA* 94:814–819.

[2] Elowitz MB, Levine AJ, Siggia ED, Swain PS (2002) Stochastic gene expression in a single cell. *Science* 297:1183–6.

[3] Colman-Lerner A, et al. (2005) Regulated cell-to-cell variation in a cell-fate decision system. *Nature* 437:699–706.

[4] Raser JM, O'Shea EK (2004) Control of stochasticity in eukaryotic gene expression. *Science* 304:1811–1814.

[5] Hilfinger A, Paulsson J (2011) Separating intrinsic from extrinsic fluctuations in dynamic biological systems. *Proc Natl Acad Sci USA* 108:12167–12172.

[6] Swain PS, Elowitz MB, Siggia ED (2002) Intrinsic and extrinsic contributions to stochasticity in gene expression. *Proceedings of the National Academy of Sciences* 99:12795–12800.

[7] Ruess J, Milias-Argeitis A, Lygeros J (2013) Designing experiments to understand the variability in biochemical reaction networks. *Journal of The Royal Society Interface* 10.

[8] Zechner C, et al. (2012) Moment-based inference predicts bimodality in transient gene expression. *Proc Natl Acad Sci USA* 109:8340–8345.

[9] Zechner C, Unger M, Pelet S, M. P, Koeppl H (2014) Scalable inference of heterogeneous reaction kinetics from pooled single-cell recordings. *Nat Methods* 11:197–202.

[10] Finkenstädt B, et al. (2013) Quantifying intrinsic and extrinsic noise in gene transcription using the linear noise approximation: An application to single cell data. *The Annals of Applied Statistics* 7:1960–1982.

[11] Hasenauer J, et al. (2011) Identification of models of heterogeneous cell populations from population snapshot data. *BMC Bioinformatics* 12:125.

[12] Shahrezaei V, Ollivier JF, Swain PS (2008) Colored extrinsic fluctuations and stochastic gene expression. *Mol Syst Biol* 4:196.

[13] N. Johnson SK (1977) *Urn Models and Their Application* (Wiley & Sons, New York).

[14] Koeppl H, Zechner C, Ganguly A, Pelet S, Peter M (2012) Accounting for extrinsic variability in the estimation of stochastic rate constants. *Int J Robust Nonlin* 22:1103–1119.

[15] Rao CV, Arkin AP (2003) Stochastic chemical kinetics and the quasi-steady-state assumption: Application to the gillespie algorithm. *The Journal of Chemical Physics* 118:4999–5010.

[16] Aalen OO, Borgan, Gjessing HK (2008) *Survival and event history analysis: a process point of view* (Springer Verlag).

[17] Bain A, Crisan D (2009) *Fundamentals of stochastic filtering* (Springer, New York).

[18] Anderson DF, Kurtz TG (2011) in *Design and Analysis of Biomolecular Circuits* (Springer), pp 3–42.

[19] Snyder DL, Miller MI (1975) *Random Point Processes in Time and Space* (Wiley & Sons, New York).

[20] Elliott RJ, Malcolm WP (2005) General smoothing formulas for markov-modulated poisson observations. *Automatic Control, IEEE Transactions on* 50:1123–1134.

[21] Anderson DF (2007) A modified next reaction method for simulating chemical systems with time dependent propensities and delays. *J Chem Phys* 127:214107.

[22] Lewis PAW, Shedler GS (1979) Simulation of nonhomogeneous Poisson processes by thinning. *Naval Research Logistics Quarterly* 26:403–413.

[23] Rosenfeld N, Young JW, Alon U, Swain PS, Elowitz MB (2005) Gene regulation at the single-cell level. *Science* 307:1962–1965.

[24] Volfson D, et al. (2006) Origins of extrinsic variability in eukaryotic gene expression. *Nature* 439:861–864.

[25] Zechner C, Deb S, Koeppl H (2013) Marginal dynamics of stochastic biochemical networks in random environments. *2013 European Control Conference (ECC)* pp 4269–4274.

[26] Bowsher CG, Voliotis M, Swain PS (2013) The fidelity of dynamic signaling by noisy biomolecular networks. *Plos Comp Biol* 9:e1002965.

[27] Jayanthi S, Nilgiriwala KS, Del Vecchio D (2013) Retroactivity controls the temporal dynamics of gene transcription. *ACS Synthetic Biology* 2:431–441.

[28] Lestas I, Vinnicombe G, Paulsson J (2010) Fundamental limits on the suppression of molecular fluctuations. *Nature* 467:174–178.